Title: **Evolving forest fire burn severity classification algorithms for multi-spectral imagery**

Author(s): Steven P. Brumby, Neal R. Harvey, Jeffrey J. Bloch, James Theiler, Simon Perkins, A. Cody Young, and John J. Szymanski

Submitted to:

<http://lib-www.lanl.gov/la-pubs/00367083.pdf>

Evolving forest fire burn severity classification algorithms for multi-spectral imagery

Steven P. Brumby, Neal R. Harvey, Jeffrey J. Bloch, James Theiler,
Simon Perkins, A. Cody Young, and John J. Szymanski

Space and Remote Sensing Sciences, Los Alamos National Laboratory,
Mail Stop D436, Los Alamos, New Mexico 87545, U.S.A.*

ABSTRACT

Between May 6 and May 18, 2000, the Cerro Grande/Los Alamos wildfire burned approximately 43,000 acres (17,500 ha) and 235 residences in the town of Los Alamos, NM. Initial estimates of forest damage included 17,000 acres (6,900 ha) of 70-100% tree mortality. Restoration efforts following the fire were complicated by the large scale of the fire, and by the presence of extensive natural and man-made hazards. These conditions forced a reliance on remote sensing techniques for mapping and classifying the burn region. During and after the fire, remote-sensing data was acquired from a variety of aircraft-based and satellite-based sensors, including Landsat 7. We now report on the application of a machine learning technique, implemented in a software package called GENIE, to the classification of forest fire burn severity using Landsat 7 ETM+ multispectral imagery. The details of this automatic classification are compared to the manually produced burn classification, which was derived from field observations and manual interpretation of high-resolution aerial color/infrared photography.

Keywords: Multispectral imagery, Genetic programming, Supervised classification, Forest fire, Wildfire.

1. INTRODUCTION: REMOTE SENSING OF FOREST FIRES

Between May 6 and May 18, 2000, the Cerro Grande/Los Alamos wildfire burned approximately 43,000 acres (17,500 ha) of forest and 235 residences in the town of Los Alamos, New Mexico (USA). Initial estimates of forest damage included 17,000 acres (6,900 ha) suffering 70-100% tree mortality. Some of the affected agencies and tribes included the United States Forest Service, the Department of Energy, the National Park Service, Santa Clara Pueblo, and the Pueblo of San Ildefonso. Restoration efforts following the fire were complicated by the large scale of the fire, and by the presence of extensive natural and man-made hazards. These conditions forced a reliance on remote sensing techniques for mapping and classifying the burn region. During and after the fire, remote-sensing data was acquired from a variety of aircraft-based and satellite-based sensors, including Landsat 7, to evaluate the impact of the fire.

Remote sensing of forest fires has traditionally involved human interpretation of visible wavelength and/or infrared photography. Since the introduction of aircraft and satellite mounted multi-spectral imaging instruments, e.g., the Advanced Very High Resolution Radiometer¹ (AVHRR) on the NOAA Polar-orbiting Operational Environmental Satellite (POES) series, and the Thematic Mapper (TM) and Enhanced Thematic Mapper (ETM+) instruments on the Landsat² series of Earth observation satellites, several physics-based and empirical algorithms for detecting forest fires have appeared in the literature. Two general approaches exist: detection of “hot-spots” and fire fronts, using, e.g., thresholds on brightness temperature^{3,4,5,6,7} in AVHRR band 3 (3.7 μ m), and mapping of post-fire burn scars. Landsat 7 fortuitously captured an image of the Cerro Grande/Los Alamos wildfire in progress on May 9, 2000 (Landsat Path 33, Row 35), in which fire fronts due to the wildfire and the back-burning efforts of the fire fighters are clearly visible (Fig. 1).

For the present work, however, we are interested in mapping and classifying the post-fire burn scar. A number of researchers have investigated the use of Landsat TM imagery for measuring wildfire impact by mapping of the burn scar. For example, Lobo et al⁸ apply a combination of spectral image segmentation and hierarchical clustering to the mapping and analysis of

---

* Work supported by U.S. Departments of Energy and Defense. Further author information: (Send correspondence to S.P.B.)
Emails: {brumby, jbloch, harve, s.perkins, jt, cody, szymanski}@lanl.gov

Figure 1. Wildfire hotspots, May 9, 2000: Bright pixels within the white boxes are wildfire hotspots. The town and laboratory of Los Alamos lie along the right edge of this image. Pixel size: 30m, ETM+ bands 7,5,3.

Figure 2. Post-fire, July 19, 2000: Bright region in center of image is the burn scar. Los Alamos town lies against the underside of the burn scar. Topography changes from forested mountains (left) to bare mesas.

fires in Mediterranean forests. Kushla and Ripple⁹ use Landsat imagery to map forest survival following a wildfire in western Oregon (USA), and investigate linear combinations of post-fire and multi-temporal TM band ratios and differences.

We now report on the application of a machine learning technique to the classification of forest fire burn severity using Landsat 7 ETM+ multispectral imagery. The details of this automatic classification are compared to a manually produced burn classification, which was derived from field observations and human photo-interpretation of high-resolution aerial color/infrared photography.

2. MACHINE LEARNING: GENETIC ALGORITHM + SUPERVISED CLASSIFIER

GENIE^{10,11,12,13} is an evolutionary computation (EC) software system, using a genetic algorithm^{14,15,16} (GA) to assemble image-processing algorithms from a collection of low-level (“primitive”) image processing operators (e.g., edge detectors, texture measures, spectral operations, and various morphological filters). This system has been shown to be effective in looking for complex terrain features, such as, e.g., golf courses¹⁷. GENIE can sequentially extract multiple features for the same scene to produce terrain classifications¹⁸. GENIE has been described at length elsewhere (see previous references), so we will only present a brief description of the system here.

GENIE follows the classic evolutionary paradigm: a population of candidate image-processing algorithms is randomly generated, and the fitness of each individual assessed from its performance in its environment, which for our case is a user-provided training scene. After fitness has been assigned, reproduction with modification follows via the evolutionary operators of selection, crossover, and mutation, applied to the most fit members of the population. The process of fitness evaluation and reproduction with modification is iterated until some stopping condition is satisfied.

The algorithms assembled by GENIE will generally combine spatial and spectral processing, and the system was in fact designed to enable spatio-spectral image processing experimentation. Each individual chromosome in the population consists of a fixed-length string of genes. Each gene in Genie corresponds to a primitive image processing operation, and so the whole chromosome describes an algorithm consisting of a sequence of primitive image processing steps. We now briefly describe our method of providing training data, our encoding of image-processing algorithms as chromosomes for manipulation by the GA, and our method for evaluating the fitness of individuals in the population.

2.1. Training Data

The environment for the population consists of one or a number of training scenes. Each training scene contains a raw multi-spectral image data cube, together with a weight plane and a truth plane. The weight plane identifies the pixels to be used in training, and the truth plane locates the features of interest in the training data. Providing sufficient quantities of good training data is a crucial to the success of any machine learning technique. In principle, the weight and truth planes may be derived from an actual ground campaign (i.e., collected on the ground at the time the image was taken), may be the result of applying some existing algorithm, and/or may be marked-up by hand using the best judgement of an analyst looking at the data. We have developed a graphical user interface (GUI), called Aladdin, for manual marking up of raw imagery. Using Aladdin, the analyst can view a multi-spectra image in a variety of ways, and can mark up training data by painting directly on the image using the mouse. Training data is ternary-valued, with the possible values being “true”, “false”, and “unknown”. True defines areas where the analyst is confident that the feature of interest does exist. False defines areas where the analyst is confident that the feature of interest does not exist. Unknown pixels do not influence the fitness of a candidate algorithm.

2.2. Representation of Image-Processing Algorithms

Traditional genetic programming¹⁹ (GP) uses a variable sized (within limits) tree representation for algorithms. Our representation differs in that it allows for reuse of values computed by sub-trees, i.e. the resulting algorithm is a graph rather than a tree. The image processing algorithm that a given chromosome represents can be thought of as a directed acyclic graph where the non-terminal nodes are primitive image processing operations, and the terminal nodes are individual image planes extracted from the multi-spectral image used as input. Our representation also differs in that the total number of nodes is fixed (although not all of these may actually be used in the final graph), and crossover is carried out directly on the linear representation.

We have restricted our “gene pool” to a set of useful primitive image processing operators (“genes”). These include spectral, spatial, logical and thresholding operators. The set of morphological operators is restricted to function-set processing morphological operators, i.e., gray-scale morphological operators having a flat structuring element. The sizes and shapes of the structuring elements used by these operators is also restricted to a pre-defined set of primitive shapes, which includes the square, circle, diamond, horizontal cross and diagonal cross, and horizontal, diagonal, and vertical lines. The shape and size of the structuring element are defined by operator parameters. Other local neighborhood/windowing operators such as mean, median, etc., specify their kernels/windows in a similar way. The spectral operators have been chosen to permit weighted sums, differences and ratios of data and/or “scratch” planes, where a scratch plane is a block of memory for storing intermediate calculations within a candidate image-processing algorithm.

A single gene consists of an operator, plus a variable number of input arguments specifying from where input is read, output arguments specifying where output is to be written, and any additional parameters that might be required to specify how the specific operator works (e.g., the diameter and shape of a structuring element used in a morphological filter). The operators used in Genie take one or more distinct image planes as input, and generally produce a single image plane as output. Input can be taken from any data plane in the training data image cube. Output is written to one of a number of scratch planes, temporary workspaces where an image plane can be stored. Genes can also take input from scratch planes, but only if that scratch plane has been written to by another gene positioned earlier in the chromosome sequence. We use a notation for genes¹⁰ that is most easily illustrated by an example: the gene [ADDP rD0 rS1 wS2] applies pixel-by-pixel addition to two input planes, read from data plane 0 and from scratch plane 1, and writes its output to scratch plane 2. Any additional required operator parameters are listed after the output arguments.

Note that although all chromosomes have the same fixed number of genes, the effective length of the resulting algorithm may be smaller than this. For instance, an operator may write to a scratch plane that is then overwritten by another gene before anything reads from it. GENIE performs an analysis of chromosome graphs when they are created and only carries out those processing steps that actually affect the final result. Therefore, the fixed length of the chromosome acts as a maximum effective length.

2.3. Supervised Classification and Fitness Evaluation

Each candidate image-processing algorithm generates a number of intermediate feature planes (or “signature” planes), which are then combined to generate a Boolean-valued mask for the feature of interest. This combination is achieved using a standard supervised classifier (we use the Fisher linear discriminant²⁰), and an optimal threshold function.

Complete (or “hard”) classification requires that the image-processing algorithm produce a binary-valued output plane for any given scene. It is possible to treat, e.g., the contents of the first scratch plane as the final output for that candidate image-

processing algorithm (thresholding would generally be required to obtain a binary result, though Genie can choose to apply its own Boolean thresholding functions). However, we have found it to be useful to perform the combination of the data and scratch planes using a non-evolutionary method, and have implemented a supervised classifier backend. To do this, we first select a subset of the scratch planes and data planes to be “signature” planes. For the present experiments, this subset consists of just the scratch planes. We then use the provided training data and the contents of the signature planes to derive the Fisher Discriminant, which is the linear combination of the signature planes that maximizes the mean separation in spectral terms between those pixels marked up as “true” and those pixels marked up as “false”, normalized by the total variance in the projection defined by the linear combination. The output of the discriminant-finding phase is a real-valued single-plane “answer” image. This is reduced to a binary image by exhaustive search over all the training pixels to find the threshold value that minimizes the total number of misclassifications (false positives plus false negatives) on the training data.

The fitness of a candidate solution is given by the degree of agreement between the final binary output plane and the training data. This degree of agreement is determined by the Hamming distance between the final binary output of the algorithm and the training data, with only pixels marked as true or false (as recorded in the weight plane) contributing towards the metric. The Hamming distance is then normalized so that a perfect score is 1000.

2.4. Software Implementation

GENIE can search a rich and complex feature space using its gene pool of standard primitive image processing operators, and the results of additional analyst-selected algorithms. The system employs both spectral and spatial image analysis techniques in combination, and can in principal simultaneously exploit data from different sensors (e.g., optical imagery plus multi-spectral imagery plus altimeter data or digital elevation models). The ability to combine diverse datasets requires that the data be co-registered, which requires use of some other package (e.g., RSI's ENVI²¹ or ERDAS's Imagine²² software packages).

Our genetic algorithm code has been implemented in object-oriented Perl. This provides a convenient environment for the string manipulations required by the evolutionary operations, and easy access to the underlying operating system (Linux). Chromosome fitness evaluation is the computationally intensive part of the evolutionary process, typically taking 90% of our total processing time. We currently use RSI's IDL²¹ language and image processing environment for this core processing, because of its visualization environment, and its ability to handle a diverse set of imagery formats. Within IDL, individual genes correspond to single primitive image operators, which are coded as IDL procedures. A chromosome can then be represented as an IDL batch executable. Many of our primitive operators do not exist in standard IDL, so we have developed an external library of C code called by IDL. In the present implementation, an IDL session is opened at the start of a run and communicates with the Perl code via a two-way UNIX pipe. This pipe is a low-bandwidth connection. Only the IDL session needs to access the input and training data (possibly hundreds of Megabytes), which requires a high-bandwidth connection. The Aladdin training data mark-up application was written in Java. Running on a single, fast Linux/Intel workstation, the system typically requires a few hours to evolve an image-processing algorithm. Re-application of an evolved image-processing algorithm to the same or a new image typically takes seconds to minutes.

3. TRAINING AND RESULTS

3.1. Training Data

The remotely-sensed images used in this paper are Landsat 7 ETM+ 30 meter multi-spectral data (ETM+ bands 1–5 and 7). These scenes are Level 1G radiance corrected and georeferenced standard data products obtained via the U.S. Geological Survey (USGS) EarthExplorer²³ web site. We used a post-fire Landsat scene from July 17, 2000, Path 34 and Row 35. The image displayed in Fig. 2 is a false-color image, which has then been converted to gray-scale and has had its contrast enhanced for the printing process. As we are interested in mapping burn scars, we generally view the data using a Visible/Infrared/Thermal pattern of a thermal IR band (ETM + band 7, 2.2 μ m) for the red component, a near IR band for the green component (band 5, 1.65 μ m), and a visible red band for the blue component (band 3, 0.66 μ m). A Landsat 7 Path/Row swath has an across-track field-of-view of approximately 185 km, with similar along-track length, resulting in a field-of-view of approximately 34,000 sq.km, which is much larger than needed for this study, and presents memory problems for our software if we attempt to ingest the whole scene. Hence, we spatially subset the image to a 1000 pixel x 1000 pixel region centered on the Los Alamos National Laboratory. We chose not to use the 60m thermal or 15m panchromatic data in the following analysis, as we wished to investigate evolution without the added complication of re-sampling of data.

We did not have any atmospheric measurements available for the scene, so we did not attempt to carry out any corrections for haze or atmosphere. The topography of Los Alamos is complex, consisting of a dormant volcano (the Jemez Mountains) rising to approximately 10,000 feet (3.3km), surrounded by a radiating network of mesas at 7,000 – 8,000 feet, falling off to the Rio Grande river valley at approximately 6,500 feet elevation. Traditionally, illumination effects due to complex

Figure 3. BAER Team burn-severity map over topographic map: Medium gray region marks high severity burn, pale gray region marks low severity/unburned region: <http://www.baerteam.org/cerrogrande>

Figure 4. Training Data over raw imagery: White patches mark “burn” regions. Gray patches mark “non-burn” regions. Note: this image is presented at a larger spatial scale than Figure 3.

topography can be approximately “factored out” by using band ratios, or removed using principal components analysis (see, e.g., Ref. 24). Here, we are interested in the GENIE software’s ability to derive results based on the raw imagery, and do not add any additional band ratio or band difference planes.

Our training data was based on the official Cerro Grande Burned-Area Emergency Rehabilitation (BAER) Team’s burn severity map, Figure 3, which was produced by trained observers flying over the fire, and visual inspection of high-resolution (~1 meter) aerial color/infrared photography collected during and immediately after the fire. Using this map as a guide, we marked up several regions of the Landsat image as almost certainly “burn”, and several regions as almost certainly “non-burn”, as shown in Figure 4. The BAER Team assign “burn severity” on the basis of tree mortality – low burn severity corresponds to grass fire and low tree mortality, medium severity burn classification implies crown fire and majority tree mortality (more than half of the trees in the marked region are dead), and the high severity burn classification requires that 70 – 100% of the trees are dead. The Cerro Grande wildfire tended to produce either high severity or low severity burn, with only a relatively small fraction of the burn classified as medium burn severity in the BAER Team maps. This was mostly due to the over-grown nature of the Ponderosa pine/mixed conifer forest which suffered most of the damage. Major species present include *Pinus ponderosa* (Ponderosa Pine), *Pseudotsuga menziesii* (Douglas fir), *Abies concolor* (White fir), *Populus tremuloides* (Aspen), *Juniperus monosperma* (Juniper), and *Pinus edulis* (Piñon). Note that we have also tried evolving algorithms from training data based purely on photo-interpretation of the 30m Landsat scene, and have obtained similar results. This is almost certainly due to the fact that in the case of the Cerro Grande wildfire the burn damage was sufficiently catastrophic that simple inspection of the 30m imagery allows accurate marking of “burn” and “non-burn” regions

3.2. Evolved Image-Processing Algorithm

The system was run for with a population of 50 chromosomes, each having a fixed length of 20 genes, and 3 intermediate feature (“scratch”) planes. The GA was allowed to evolve for 30 generations, in this case, evaluating 1282 distinct candidate image processing algorithms, which is very small compared to search space of possible algorithms given our representation. This required approximately 7 hours of wall-clock time running on a 500MHz Linux/Intel Pentium 2 workstation.

The best evolved image-processing algorithm had the chromosome,

```
[OPEN rD1 wS1 1 1][ADDS rD4 wS3 0.34][NEG rS1 wS1][MULTP rD4 rS3 wS2]
[LINCOMB rS1 rD6 wS3 0.11][ADDP rS1 rS3 wS1][SUBP rS1 rD5 wS1]
```

Figure 5. Evolved image-processing algorithm: each node (gene) in the graph is labeled by its position along the length of the chromosome and by the GENIE software’s mnemonic for the primitive image processing operator.

and is shown as a diagram in Figure 5, where each node (gene) in the graph is labeled by its position along the length of the chromosome and by our mnemonic for the operator (e.g., 1. OPEN for a morphological opening operator). In words, the image-processing algorithm works as follows. Note that GENIE converts the byte-valued raw data to real-valued data (64 bit doubles) and keeps that precision through all its calculations.

1. Data plane D1 (ETM+ band 1, visible blue 0.48 μ m) undergoes a grayscale morphological opening operation (node 1. OPEN) using a “circular” structuring element with diameter equal to 3 pixels (equivalent to a 3x3 square with corners removed) and the result is written to scratch plane S1,
2. The negative of this plane is taken (node 3. NEG), i.e., $S1 \rightarrow -S1$,
3. The new S1 is linearly combined (node 5. LINCOMB) with data plane D6 (ETM+ band 7, medium wavelength infrared (MWIR) 2.22 μ m) with linear weights: $0.11*S1 + 0.89*D6$ and the result written to scratch plane S3 (its final value),
4. Scratch planes S1 and S3 are summed (node 6. ADDP), and the difference (node 7. SUBP) of this sum and data plane D5 (ETM+ band 5, MWIR 1.65 μ m), $S1 + S3 - D5$, is written to S1 (its final value),
5. Data plane D4 (ETM+ band 4, near infrared 0.83 μ m) has a constant, 0.34 times a DATASCALE variable equal to the range of the input raw data values, added to each pixel (node 2. ADDS) and is multiplied by D4 again to form the linear combination $D4*D4 + (0.34*DATASCALE)*D4$, which is written to scratch plane S2 (its final value).

The final values of S1, S2, and S3 are then combined in the linear sum, where the coefficients and intercept have been chosen by the Fisher discriminant, as described in Section 2.3, above, to produce our real-valued answer plane A (Figure 6):

$$A = 0.0147*S1 - 0.0142*S2 + 0.0134*S3 + 1.554$$

The optimal threshold found by GENIE, given the training data, was 0.3437. Converting A to a Boolean mask at that threshold value produces Figure 7. In relation to the BAER map, Figure 3, we see that the system has extracted the high, medium, and low severity burn regions, but also presents a number of false positives. On inspection, these turn out to correspond to two physical categories of land cover: bare ground/rock, and cloud shadows. The histogram of A shows a bimodal distribution (Figure 8), as expected if the burn/non-burn classes are separable. Adjusting the threshold on A to fall at the between-peak minimum of the histogram at 0.7930 (a different optimization criterion for the threshold than that used by default by GENIE) produces a new Boolean mask, Figure 9, in which almost all the false positives have been removed, and the remaining pixels marked as “burn” correspond very closely to the high severity burn regions in the BAER map

Figure 6. Real-valued Answer Plane: We use a Fisher Discriminant to find the optimal linear combination of evolved “signature” planes into a real-valued answer plane. Regions which will tend to be classified as “burn” are bright. This image has been histogram-equalized to increase contrast.

Figure 7. Burn mask: GENIE determines an optimal threshold for converting the real-valued answer plane to a Boolean mask. Misidentified pixels are mostly cloud shadows (e.g., compact regions on left), or bare ground/rock (lower right and bottom).

3.3. Application to Non-Training Data

The evolved algorithm can now be applied to any scene. To check the reasonableness of our algorithm’s performance, we ran the image-processing algorithm over a larger fraction of the Landsat scene, encompassing the entire Jemez mountain range. The result is shown in Figure 10. We claim that this overall result is quite reasonable, and only fails where the Landsat swath ends (which can be easily masked out). Of particular interest is the persistent detection of a severe burn site on the Western side of the Jemez mountains, Figures 10, 11, 12, which cannot obviously be excluded due to cloud shadows or data drop-out. In fact, this turns out to be a true detection of a second wildfire, the Stable wildfire (effecting Stable Stream and School House Mesa in the Jemez Mountains of northern New Mexico), which destroyed approximately 800 acres of forest in September/October of 1999. As GENIE had no knowledge of this fire during its training, we find this detection, together with the reasonable behavior of the evolved image-processing algorithm over this large region, as quite encouraging for the future usefulness of this machine learning technique.

4. CONCLUSIONS

We have investigated evolution of an image-processing algorithm to extract wildfire burn scars in Landsat 7 ETM+ imagery, and have described the operation of the evolved algorithm in some detail. The evolved algorithm shows a good qualitative fit to the published BAER Team burn-severity map of the May 2000 Cerro Grande/Los Alamos wildfire, specifically in comparison to their high-severity burn class (70-100% tree mortality regions). The algorithm can be confused by dark cloud shadows, and by bare ground/rock outcrops which are physically very similar to the charred remains of the severely burned forest, but adjustment of its final threshold can significantly improve this behavior. Applying the algorithm outside the training area showed that it continued to produce reasonable results over a large spatial region, and in fact was able to detect a second small wildfire on the west side of the Jemez mountains (September/October 1999 Stable wildfire). We find these results quite encouraging for the future application of this machine learning technique.

Figure 8. Histogram of the Answer Plane (Fig. 6): The bimodal distribution indicates that “burn” and “non-burn” are indeed separable classes.

Figure 9. Final burn mask: Thresholding the answer plane at the between-peak minimum of the bimodal distribution produces this burn mask, which has almost eliminated false positives. There is substantial agreement with the details of the BAER map (Fig.3).

ACKNOWLEDGEMENTS

The authors would like to thank Leslie Hanson, Steven Koch, and Randy Balice of the Ecology Group for useful discussions and access to data used in this work. The GENIE system is the result of the combined efforts of a team of people at LANL, including, in addition to the authors of this paper: Reid Porter, Mark Galassi, Kevin Lacker, and Melanie Mitchell.

REFERENCES

1. The National Oceanographic and Atmospheric Administration (NOAA) POES satellites and the AVHRR instrument are described on the NOAA web site <http://www.ncdc.noaa.gov>
2. Landsat TM and ETM+ are described on the U.S. Geological Survey (USGS) web site <http://landsat7.usgs.gov>
3. Y.J. Kaufmann, C.J. Tucker, and I. Fung, “Remote sensing of biomass burning in the tropics”, *J. Geophysical Research*, Vol 95, No. D7, pp. 9927-9939, 1990, and references therein.
4. Y. Rauste, E. Herland, H. Frelander, K. Soini, T. Kuoremäki, and A. Ruokari, “Satellite-based forest fire detection for fire control in boreal forests”, *Int. J. Remote Sensing*, Vol. 18, No. 12, pp. 2641-2656, 1997.
5. N.P. Minko, N.A. Abushenko, V.V. Koshelev, “Forest fire detection in East Siberia forests using AVHRR/NOAA data”, *Proc. SPIE*, Vol. 3502, pp. 192-200, 1998.
6. R. Lasaponara, V. Cuomo, V. Tramutoli, N. Pergola, C. Pietrapertosa, and T. Simoniello, “Forest fire danger estimation based on the integration of satellite AVHRR data and topographic factors”, *Proc. SPIE*, Vol. 3868, pp. 241-252, 1999.
7. S.H. Boles and D.L. Verbyla, “Effect of scan angle on AVHRR fire detection accuracy in interior Alaska”, *Int. J. Remote Sensing*, Vol. 20, No. 17, 3437-3443, 1999.
8. A. Lobo, N. Pineda, R. Navarro-Cedillo, P. Fernandez-Rebollo, F.J. Salas, J.-L. Fernández-Turiel, and A. Fernández-Palacios, “Mapping forest fire impact from Landsat TM imagery”, *Proc. SPIE*, Vol. 3499, 340-347, 1998.
9. J.D. Kushla and W.J. Ripple, “Assessing wildfire effects with Landsat thematic mapper data”, *Int. J. Remote Sensing*, Vol. 19, No. 13, 2493-2507.
10. S.P. Brumby, J. Theiler, S.J. Perkins, N.R. Harvey, J.J. Szymanski, J.J. Bloch, and M. Mitchell, “Investigation of feature extraction by a genetic algorithm”, *Proc. SPIE*, Vol. 3812, pp. 24-31, 1999.
11. J. Theiler, N.R. Harvey, S.P. Brumby, J.J. Szymanski, S. Alferink, S.J. Perkins, R. Porter, and J.J. Bloch, “Evolving retrieval algorithms with a genetic programming scheme”, *Proc. SPIE*, Vol. 3753, pp. 416-425, 1999.

Figure 10 Testing the evolved image-processing algorithm: Extended region (top) and burn mask (bottom). The evolved image-processing algorithm continues to work well, except at the edge of the Landsat swath (bottom image, lower right). The small black region on the left of the burn mask represents a true detection of a second recent wildfire, the 1999 Stable wildfire.

Figure 11. Detail of the second detected burn: Grayscale image, ETM+ band 7.

Figure 12. Burn mask for second detected burn: Location agrees with the known 1999 Stable wildfire.

12. N.R. Harvey, S.P. Brumby, S.J. Perkins, R.B. Porter, J. Theiler, A.C. Young, J.J. Szymanski, and J.J. Bloch, "Parallel evolution of image processing tools for multispectral imagery", Proc. SPIE, Vol. 4132, pp. 72-82, 2000.
13. S. Perkins, J. Theiler, S.P. Brumby, N.R. Harvey, R.B. Porter, J.J. Szymanski, and J. J. Bloch, "GENIE: A hybrid genetic algorithm for feature classification in multi-spectral images", Proc. SPIE, Vol. 4120, pp 52-62, 2000.
14. J. H. Holland, *Adaptation in Natural and Artificial Systems*, University of Michigan, Ann Arbor (1975).
15. I. Rechenberg, *Evolutionsstrategie: Optimierung technischer Systeme nach Prinzipien der biologischen Evolution*, Fromman-Holzboog, Stuttgart (1973).
16. L. Fogel, A. Owens and M. Walsh, *Artificial Intelligence through Simulated Evolution*, Wiley, New York (1966).
17. N.R. Harvey, S. Perkins, S.P. Brumby, J. Theiler, R.B. Porter, A.C. Young, A.K. Varghese, J.J. Szymanski, and J.J. Bloch, "Finding golf courses: The ultra high tech approach", Proc. Second European Workshop on Evolutionary Computation in Image Analysis and Signal Processing (EvoIASP2000), Edinburgh, UK, pp 54-64, 2000.
18. S.P. Brumby, N.R. Harvey, S. Perkins, R.B. Porter, J.J. Szymanski, J. Theiler, and J.J. Bloch, "A genetic algorithm for combining new and existing image processing tools for multispectral imagery", Proc. SPIE, Vol. 4049, pp. 480-490, 2000.
19. J. R. Koza, *Genetic Programming: On the Programming of Computers by Natural Selection*, MIT, Cambridge (1992).
20. For example, see C.M. Bishop, *Neural Networks for Pattern Recognition*, pp.105–112, Oxford University (1995).
21. For details on RSI Inc.'s ENVI and IDL software, see <http://www.rsinc.com>
22. For details on ERDAS Inc.'s Imagine software, see <http://www.erdas.com>
23. For details on USGS EarthExplorer, see <http://edcsns17.cr.usgs.gov/EarthExplorer>.
24. R.A. Schowengerdt, *Remote Sensing*, 2nd ed., Academic, San Diego (1997).